

Nonparametric Nearest Neighbor Random Process Clustering

Michael Tschannen and Helmut Bölcskei

Dept. IT & EE, ETH Zurich, Switzerland

Email: {michaelt, boelcskei}@nari.ee.ethz.ch

Abstract—We consider the problem of clustering noisy finite-length observations of stationary ergodic random processes according to their nonparametric generative models without prior knowledge of the model statistics and the number of generative models. Two algorithms, both using the L^1 -distance between estimated power spectral densities (PSDs) as a measure of dissimilarity, are analyzed. The first algorithm, termed nearest neighbor process clustering (NNPC), to the best of our knowledge, is new and relies on partitioning the nearest neighbor graph of the observations via spectral clustering. The second algorithm, simply referred to as k -means (KM), consists of a single k -means iteration with farthest point initialization and was considered before in the literature, albeit with a different measure of dissimilarity and with asymptotic performance results only. We show that both NNPC and KM succeed with high probability under noise and even when the generative process PSDs overlap significantly, all provided that the observation length is sufficiently large. Our results quantify the tradeoff between the overlap of the generative process PSDs, the noise variance, and the observation length. Finally, we present numerical performance results for synthetic and real data.

I. INTRODUCTION

Consider a set of N noisy length- M observations of stationary ergodic random processes stemming from $L < N$ (typically $L \ll N$) different generative models. We want to cluster these observations according to their generative models without prior knowledge of the model statistics and the number of generative models L . This problem arises in many domains of science and engineering where large amounts of data have to be divided into meaningful subsets in an unsupervised fashion, e.g., for efficient processing or storage. Examples include clustering of audio and video sequences [1], electrocardiography (ECG) recordings [2], and industrial production indices [3].

Many existing random process clustering methods quantify the dissimilarity between observations using the Euclidean distance between either estimated model parameters [3], or estimated cepstral coefficients [2], [4], or normalized periodograms [5]. Other methods rely on divergences (e.g., Kullback-Leibler divergence) between normalized periodograms [6], [7] or use the distributional distance between processes [8], [9]. In all cases the resulting distances are then fed to a standard clustering algorithm such as k -means or hierarchical clustering. Another line of work employs a Bayesian framework to infer the cluster assignments, e.g., according to a maximum a posteriori criterion [10]. While many of these approaches have proven effective in practice, corresponding analytical performance results are scarce. Moreover, existing

analytical results are mostly concerned with the asymptotic regime where the observation length tends to infinity (see, e.g., [3], [6]–[9]); the finite observation-length regime has attracted virtually no attention.

Contributions: We consider two process clustering algorithms that apply to nonparametric generative models and measure dissimilarity between observations through the L^1 -distance between estimated power spectral densities (PSDs). The first algorithm, termed nearest neighbor process clustering (NNPC), is inspired by the subspace clustering algorithm described in [11] and, to the best of our knowledge, has not been considered in the literature before. NNPC relies on partitioning the q -nearest neighbor graph of the observations via spectral clustering (the number of nearest neighbors q is a parameter of NNPC). The second algorithm, which will simply be referred to as k -means (KM), consists of a single k -means iteration with farthest point initialization [12] and was first proposed in [8] with a different distance measure. The original formulation of KM was shown in [8] to deliver the correct segmentation with probability tending to 1 as the observation length $M \rightarrow \infty$.

Assuming real-valued zero-mean stationary ergodic Gaussian processes as generative models, we characterize the performance of NNPC and KM analytically for finite-length observations contaminated by independent additive white Gaussian noise. We find that both algorithms succeed with high probability even when the (true) PSDs exhibit significant overlap, all provided that the observation length is sufficiently large. Our analytical results quantify the tradeoff between observation length, noise variance, and distance between the (true) PSDs of the generative models. Finally, we evaluate the performance of the two algorithms on synthetic and on real data.

Notation: We use lowercase boldface letters to denote vectors, uppercase boldface letters to designate matrices, and the superscript T stands for transposition. v_i is the i th entry of the vector \mathbf{v} . For the matrix \mathbf{A} , $\mathbf{A}_{i,j}$ designates the entry in row i and column j , $\|\mathbf{A}\|_{2 \rightarrow 2}$ its maximum singular value, and $\|\mathbf{A}\|_F := (\sum_{i,j} |\mathbf{A}_{i,j}|^2)^{1/2}$ its Frobenius norm. \mathbf{I} stands for the identity matrix. The i th element of a sequence x is denoted by $x[i]$. For a positive integer N , $[N]$ designates the set $\{1, 2, \dots, N\}$. $\mathbb{E}[X]$ is the expectation of the random variable X and the notation $Y \sim X$ indicates that the random variable Y has the same distribution as X . We say that a subgraph H of a graph G is connected if every pair of nodes

in H can be joined by a path with nodes exclusively in H . A connected subgraph H of G is called a connected component of G if there are no edges between H and the remaining nodes in G [13].

II. FORMAL PROBLEM STATEMENT AND ALGORITHMS

We consider the following clustering problem: Given the unlabeled data set $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_L$ of cardinality N , where $\mathcal{X}_\ell = \{\hat{x}_i^{(\ell)}\}_{i=1}^{n_\ell}$ contains the contiguous noisy length- M observations $\hat{x}_i^{(\ell)}$ of the real-valued stationary ergodic random process $X^{(\ell)}[m]$, $m \in \mathbb{Z}$, corresponding to the ℓ th generative model, find the partition $\mathcal{X}_1, \dots, \mathcal{X}_L$. The statistics of the generative models and of the noise processes, as well as the number of generative models, are all assumed unknown.

Both clustering algorithms we consider are based on the following distance measure. Denoting the PSD of $X^{(\ell)}$ by $s^{(\ell)}(f)$, $f \in [0, 1]$, we define the distance between $X^{(k)}$ and $X^{(\ell)}$ by $d(X^{(k)}, X^{(\ell)}) := \frac{1}{2} \int_0^1 |s^{(k)}(f) - s^{(\ell)}(f)| df$. For processes $X^{(\ell)}$ with unit power, i.e., $\int_0^1 s^{(\ell)}(f) df = 1$, $\ell \in [L]$, the factor $1/2$ ensures that $d(X^{(k)}, X^{(\ell)})$ takes on values in $[0, 1]$. $d(X^{(k)}, X^{(\ell)})$ essentially quantifies the dissimilarity between the support sets of $s^{(k)}$ and $s^{(\ell)}$, i.e., it is close to 1 when $s^{(k)}$ and $s^{(\ell)}$ are concentrated on disjoint frequency bands.

We now present the NNPC and the KM algorithms. Recall that NNPC is inspired by the subspace clustering algorithm introduced in [11], and KM is obtained by replacing the distance measure in Algorithm 1 in [8] by the distance measure d defined above.

The NNPC algorithm. *Given a set \mathcal{X} of N length- M observations, the number of generative models L (the estimation of L from \mathcal{X} is discussed below), and the parameter q , carry out the following steps:*

Step 1: *For every $\hat{x}_i \in \mathcal{X}$, estimate the PSD $\hat{s}_i(f)$ via the Blackman-Tukey (BT) estimator according to*

$$\hat{s}_i(f) := \sum_{m=-M+1}^{M-1} g[m] \hat{r}_i[m] e^{-i2\pi f m}, \quad \text{where} \quad (1)$$

$$\hat{r}_i[m] := \frac{1}{M} \sum_{n=1}^{M-|m|} \hat{x}_i[n+m] \hat{x}_i[n], \quad |m| \leq M-1,$$

and $g[m]$, $m \in \mathbb{Z}$, is an even window function (i.e., $g[m] = g[-m]$) with $g[0] = 1$ and nonnegative bounded discrete-time Fourier transform $g(f)$, i.e., $0 \leq g(f) \leq A < \infty$, $f \in [0, 1]$. Identify the set $\mathcal{T}_i \subset [N] \setminus \{i\}$ of cardinality q defined through

$$d(\hat{x}_i, \hat{x}_j) \leq d(\hat{x}_i, \hat{x}_p), \quad \text{for all } j \in \mathcal{T}_i \text{ and all } p \notin \mathcal{T}_i.$$

Step 2: *Let $\mathbf{z}_j \in \mathbb{R}^N$ be the vector with i th entry $\exp(-2d(\hat{x}_i, \hat{x}_j))$, if $j \in \mathcal{T}_i$, and 0, if $j \notin \mathcal{T}_i$.*

Step 3: *Construct the adjacency matrix \mathbf{A} according to $\mathbf{A} = \mathbf{Z} + \mathbf{Z}^T$, where $\mathbf{Z} = [\mathbf{z}_1 \dots \mathbf{z}_N]$.*

Step 4: *Apply normalized spectral clustering [13] to (\mathbf{A}, L) .*

The KM algorithm [8]. *Given a set \mathcal{X} of N length- M observations and the number of generative models L , carry out the following steps:*

Step 1: *Initialize $c_1 := 1$ and $\hat{\mathcal{X}}_\ell := \{\}$ for all $\ell \in [L]$.*

Step 2: *For every $\hat{x}_i \in \mathcal{X}$, estimate the PSD $\hat{s}_i(f)$ via the BT*

estimator (1).

Step 3: *for $p = 2$ to L do:*

$$c_p := \arg \max_{i \in [N]} \min_{\ell \in [p-1]} d(\hat{x}_i, \hat{x}_{c_\ell}).$$

Step 4: *for $i = 1$ to N do:*

$$\ell^* \leftarrow \arg \min_{\ell \in [L]} d(\hat{x}_i, \hat{x}_{c_\ell})$$

$$\hat{\mathcal{X}}_{\ell^*} \leftarrow \hat{\mathcal{X}}_{\ell^*} \cup \{\hat{x}_i\}$$

Both algorithms are based on comparing distances between observations, and are therefore meaningful only if the observations are of comparable average power.

We henceforth denote the nearest neighbor graph with adjacency matrix \mathbf{A} obtained in Step 3 of NNPC by G . The parameter q in the NNPC algorithm determines the number of edges in G . Choosing q too small may result in the observations stemming from a given generative model forming multiple connected components in G and hence not being assigned to the same cluster. This problem can be countered by increasing q , which, however, increases the chances of observations coming from different generative models being connected in G and hence being misclustered. The issue of how to choose q in practice is further discussed in the next section.

In the NNPC algorithm L may be estimated using the *eigengap heuristic* [13], which relies on the fact that the number of zero eigenvalues of the normalized Laplacian of G corresponds to the number of connected components of G .

We finally note that the BT PSD estimates (1) can be computed efficiently via the FFT.

III. PERFORMANCE RESULTS

For our analytical performance results, we assume that the $\hat{x}_i^{(\ell)}$, for given ℓ , are obtained as contiguous length- M observations of $\tilde{X}^{(\ell)}[m] := X^{(\ell)}[m] + W^{(\ell)}[m]$, $m \in \mathbb{Z}$, where $X^{(\ell)}$ is zero-mean stationary Gaussian with PSD $s^{(\ell)}(f)$, and $W^{(\ell)}$ is a zero-mean white Gaussian noise process with variance σ^2 and independent of $X^{(\ell)}$. Note that the noise statistics are identical across the generative models. The autocorrelation functions (ACFs) $r^{(\ell)}[m] := \int_0^1 s^{(\ell)}(f) e^{i2\pi f m} df$ are assumed absolutely summable, i.e., $\sum_{m=-\infty}^{\infty} |r^{(\ell)}[m]| < \infty$, $\ell \in [L]$, which implies ergodicity of the corresponding $X^{(\ell)}$. Further, we assume that the PSDs are normalized according to $\int_0^1 s^{(\ell)}(f) df = 1$, $\ell \in [L]$, and we define $B := \max_{\ell \in [L]} \sup_{f \in [0, 1]} s^{(\ell)}(f)$. Our performance results depend on the maximum ACF moment $\mu_{\max} := \max_{\ell \in [L]} \mu^{(\ell)}$, where $\mu^{(\ell)} := \sum_{m=-\infty}^{\infty} |h[m]| |r^{(\ell)}[m]|$ with

$$h(m) := \begin{cases} 1 - g(m)(1 - |m|/M), & \text{for } |m| < M \\ 1, & \text{otherwise.} \end{cases} \quad (2)$$

For each ℓ , the $\hat{x}_i^{(\ell)}$ may either stem from independent realizations of $\tilde{X}^{(\ell)}$ or correspond to different (possibly overlapping) length- M segments of a given realization of $\tilde{X}^{(\ell)}$. In the latter case the $\hat{x}_i^{(\ell)}$ will not be statistically independent in general.

Our main result for NNPC ensures the *no false connections (NFC) property* defined as follows.

Definition 1 (No False Connections Property). *G has no false connections if, for all $\ell \in [L]$, nodes corresponding to \mathcal{X}_ℓ are connected to other nodes corresponding to \mathcal{X}_ℓ only.*

Although the NFC property alone does not guarantee correct clustering, it was found to be a sensible performance measure for subspace clustering algorithms (see, e.g., [11], [14]). To ensure correct clustering one would additionally have to ensure that the subgraphs of G corresponding to the \mathcal{X}_ℓ are connected. Proving connectivity for NNPC appears to be difficult in the noisy finite observation-length regime.

Theorem 1. *Let $q \leq \min_{\ell \in [L]}(n_\ell - 1)$ and let \mathcal{X} be generated according to the data model described above. Then, the clustering condition*

$$\min_{\substack{k, \ell \in [L]: \\ k \neq \ell}} d(X^{(k)}, X^{(\ell)}) > 8A(B + \sigma^2) \sqrt{\frac{2 \log M}{M}} + 2\mu_{\max} \quad (3)$$

guarantees that G has NFC with probability at least $1 - 2N/M^2$.

Our main result for KM comes in terms of a stronger performance guarantee, namely it ensures correct clustering. This is thanks to the fact that KM does not have a spectral clustering step and is hence much easier to analyze. On the other hand NNPC typically outperforms KM in practice.

Theorem 2. *Let \mathcal{X} be generated according to the data model described above. Then, under the clustering condition (3), the partition $\hat{\mathcal{X}}_1, \dots, \hat{\mathcal{X}}_L$ of \mathcal{X} inferred by KM corresponds to the true partition $\mathcal{X}_1, \dots, \mathcal{X}_L$ with probability at least $1 - 2N/M^2$.*

The proofs of Theorems 1 and 2 are given in the Appendix. Theorems 1 and 2 essentially state that NNPC and KM succeed even when the PSDs $s^{(\ell)}$ of the $X^{(\ell)}$ overlap significantly and the observations are contaminated by strong noise, all this provided that the observation length M is sufficiently large and the window g smoothes the BT PSD estimates appropriately so that μ_{\max} is small. The clustering condition (3) quantifies the tradeoff between the (maximum) amount of overlap of the $s^{(\ell)}$ (through $\min_{k, \ell \in [L]: k \neq \ell} d(X^{(k)}, X^{(\ell)})$), the observation length M , and the noise variance σ^2 . The first term on the RHS of (3) vanishes as $M \rightarrow \infty$. For $r^{(\ell)}$ with small essential support relative to M , i.e., $r^{(\ell)}[m] \approx 0$ for $m \geq M_0$ with $M_0 \ll M$, choosing g such that $g[m] \approx 1$ for $m \leq M_0$ yields $\mu_{\max} \ll 1$ (since $h[m] \approx 1 - g[m] \approx 0$ for $m \leq M_0$). Hence, the clustering condition (3) can indeed be satisfied if the $r^{(\ell)}$ decay rapidly enough. Note that the choice of g affects the constant A . In particular, increasing the width of g results in larger A .

To ensure that the probability of success of NNPC and KM is high, we need to take $M \gg \sqrt{N}$, i.e., the observation length should be large relative to the square root of the number of observations.

We next address the choice of q for NNPC. The condition $q \leq \min_{\ell \in [L]}(n_\ell - 1)$ in Theorem 1 depends on the unknown quantities n_ℓ and admits a large range of values for q only if the clusters have balanced sizes, i.e., if $n_\ell \approx N/L$, $\ell \in [L]$. In practice, however, the performance of NNPC is observed to be quite robust w.r.t. the choice of q even when the cluster sizes are imbalanced.

We would like to point out that virtually all existing analytical performance results for random process clustering apply to the asymptotic regime where $M \rightarrow \infty$, for N

fixed. The results that come closest to ours in spirit can be found in [6], [7], where it is shown, in the asymptotic setting, that observations coming from different generative models can consistently (in the statistical sense) be discriminated via a PSD-based distance measure, provided that the PSDs of the generative models differ on a set of positive Lebesgue measure.

Finally, we note that the random process clustering problem considered here can be cast as a subspace clustering problem simply by interpreting the observations \hat{x}_i as vectors in \mathbb{R}^M . Numerical results, not reported here, demonstrate, however, that NNPC clearly outperforms its subspace clustering cousin, the thresholding based subspace clustering (TSC) algorithm [11]. This is thanks to NNPC exploiting the stationarity of the generative models, a property that is usually inexistent in subspace clustering and is hence not taken into account by TSC.

IV. NUMERICAL RESULTS¹

Synthetic data: We evaluate the performance of NNPC and KM in terms of the clustering error (CE), i.e., the fraction of misclustered observations. We consider $L = 3$ generative ARMA models with PSDs $s_{\mathbf{a}, \mathbf{b}}(f) = |\sum_{u=1}^{m+1} b_u e^{-i2\pi u f}|^2 / |\sum_{v=1}^{n+1} a_v e^{-i2\pi v f}|^2$, where \mathbf{a} and \mathbf{b} are coefficient vectors of length $n+1$ and $m+1$, respectively. We choose $\mathbf{a}_1 = 1$, $\mathbf{b}_1 = [3/4 \ 1 \ -7/4 \ 1/2]$, $\mathbf{a}_2 = 1$, $\mathbf{b}_2 = [1/2 \ 5/4 \ -3/2 \ 3/4]$, and $\mathbf{a}_3 = [1 \ -1/5 \ 2/5 \ 1/10]$, $\mathbf{b}_3 = 1$, and then normalize the coefficient vectors to ensure that the processes have unit power. It can be seen in Figure 1 (left) that the process PSDs overlap significantly. For the BT PSD estimator, we use a Gaussian window g with standard deviation 50 for both NNPC and KM. We sample $n_\ell = 25$ independent observations from each generative model, and for each value of M the CE is averaged over 200 such problem instances. The number of generative models $L = 3$ is assumed known. The performance of NNPC is rather insensitive to the choice of q for $10 \leq q \leq 20$ (corresponding results are not shown here) and we set $q = 10$. Figure 1 (right) shows the resulting average CE as a function of M . NNPC is seen to consistently outperform KM. For larger M the performance difference between NNPC and KM becomes less pronounced. Note that for the choice of model parameters and window function in this experiment the clustering condition (3) is not satisfied.




Fig. 1. PSDs of the generative models (left) and average CE as a function of M (right).

Real data: We consider the problem of clustering sequences of human locomotion according to the activities performed. Specifically, we repeat the experiment from [9], [15],

¹Matlab code available at <http://www.nari.ee.ethz.ch/commth/research/>

which uses the Carnegie Mellon Motion Capture database² containing motion sequences of 149 subjects performing various activities. The motion sequences describe the positions of markers on different body parts over time, recorded using optical tracking. The experiment is based on subjects #16 and #35 for which the database contains 49 and 33 sequences, respectively, labeled either as “walking” or “running”. We assume $L = 2$ to be known and we cluster the sequences describing the motion of the marker placed on the right foot of the subject. Differences in sequence lengths are accounted for by zero-padding to the maximum sequence length. Furthermore, we normalize the BT PSD estimates to unit power. Table I lists the CE as well as the entropy of the clustering confusion matrix S (defined in [15, Sec. 6]) for $q = 6$ (i.e., the value in $[\min_{\ell \in [L]}(n_\ell - 1)]$ which yields the lowest CE and S concurrently). Comparing S to the corresponding values reported in [9], [15], we find that for subject #35 NNPC and KM perform better than the algorithm proposed in [15] and match the performance of the algorithm considered in [9], while for subject #16 NNPC outperforms both of these algorithms as well as KM.

subject	NNPC		KM	
	CE	S	CE	S
#16	0.02	0.09	0.24	0.55
#35	0	0	0	0

TABLE I
CE AND S FOR CLUSTERING HUMAN MOTION SEQUENCES

APPENDIX

The central element in the proofs of Theorems 1 and 2 is the following result, proven at the end of the Appendix.

Theorem 3. *Consider a data set \mathcal{X} generated according to the data model described in Section III. Then, the clustering condition (3) implies*

$$\min_{\substack{k, \ell \in [L]: \\ k \neq \ell}} \min_{\substack{i \in [n_\ell]: \\ j \in [n_k]}} d(\hat{x}_j^{(k)}, \hat{x}_i^{(\ell)}) > \max_{\ell \in [L]} \max_{\substack{i, j \in [n_\ell]: \\ i \neq j}} d(\hat{x}_i^{(\ell)}, \hat{x}_j^{(\ell)}) \quad (4)$$

with probability at least $1 - 2N/M^2$.

Theorem 3 essentially says that under the clustering condition (3) observations stemming from the same generative model are closer (in terms of the distance measure d) than observations stemming from different generative models. We now show how Theorems 1 and 2 follow from Theorem 3.

Proof of Theorem 1: The NFC property for $q \leq \min_{\ell \in [L]}(n_\ell - 1)$ is a direct consequence of (4), which by Theorem 3, is implied by the clustering condition (3).

Proof of Theorem 2: The proof is effected by showing that in Step 3 KM selects an observation with a different generative model in every iteration, i.e., $\{\hat{x}_{c_\ell}\}_{\ell=1}^L$ contains exactly one observation of each generative model, provided that the clustering condition (3) and hence, by Theorem 3, (4) holds. The argument is then concluded by noting that, again by (4) and hence by (3), the partition $\hat{\mathcal{X}}_1, \dots, \hat{\mathcal{X}}_L$ obtained in Step 4 corresponds to the true partition $\mathcal{X}_1, \dots, \mathcal{X}_L$.

Suppose that after the p th iteration in Step 3 of KM the observations $\hat{x}_{c_\ell}, \ell \in [p]$, all come from different generative models, and assume w.l.o.g. that the generative model underlying \hat{x}_{c_ℓ} has index ℓ , for $\ell \in [p]$. For iteration $p+1$, it follows from (4) that

$$\begin{aligned} & \max_{i \in [N]} \min_{\ell \in [p]} d(\hat{x}_i, \hat{x}_{c_\ell}) \\ &= \max \left\{ \max_{\substack{k \in [p], \\ i \in [n_k]}} \min_{\ell \in [p]} d(\hat{x}_i^{(k)}, \hat{x}_{c_\ell}^{(\ell)}), \max_{\substack{k \in [L] \setminus [p], \\ i \in [n_k]}} \min_{\ell \in [p]} d(\hat{x}_i^{(k)}, \hat{x}_{c_\ell}^{(\ell)}) \right\} \\ &= \max \left\{ \underbrace{\max_{\substack{k \in [p], \\ i \in [n_k]}} d(\hat{x}_i^{(k)}, \hat{x}_{c_k}^{(k)})}_{\leq \max_{\ell \in [L]} \max_{\substack{i, j \in [n_\ell]: \\ i \neq j}} d(\hat{x}_i^{(\ell)}, \hat{x}_j^{(\ell)})}, \underbrace{\max_{\substack{k \in [L] \setminus [p], \\ i \in [n_k]}} \min_{\ell \in [p]} d(\hat{x}_i^{(k)}, \hat{x}_{c_\ell}^{(\ell)})}_{\geq \min_{\substack{k, \ell \in [L]: \\ k \neq \ell}} \min_{\substack{i \in [n_\ell], \\ j \in [n_k]}} d(\hat{x}_i^{(k)}, \hat{x}_j^{(\ell)})} \right\} \\ &= \max_{\substack{k \in [L] \setminus [p], \\ i \in [n_k]}} \min_{\ell \in [p]} d(\hat{x}_i^{(k)}, \hat{x}_{c_\ell}^{(\ell)}). \end{aligned} \quad (5)$$

Note that in the maximization in (5) k runs over $[L] \setminus [p]$ which means that $\hat{x}_{c_{p+1}}$ is guaranteed to have a generative model that is different from those underlying $\hat{x}_{c_1}, \dots, \hat{x}_{c_p}$. Iterating the preceding argument for $p = 2, \dots, L$, we find that, indeed, $\{\hat{x}_{c_\ell}\}_{\ell=1}^L$ contains exactly one observation of each generative model.

Proof of Theorem 3: Let $\tilde{s}^{(\ell)}(f) := s^{(\ell)}(f) + w^{(\ell)}(f)$ with $w^{(\ell)}(f) = \sigma^2$, $f \in [0, 1]$, for all $\ell \in [L]$, be the PSD of $\tilde{X}^{(\ell)}$ and $\tilde{r}^{(\ell)}$ the corresponding ACF. Define $e_i^{(\ell)}(f) := \hat{s}_i^{(\ell)}(f) - \tilde{s}^{(\ell)}(f)$ and set $\varepsilon := \max_{\ell \in [L]} \sup_{f \in [0, 1]} |e_i^{(\ell)}(f)|$. We have for all $k, \ell \in [L], i \in [n_\ell], j \in [n_k]$, that

$$\begin{aligned} d(\hat{x}_j^{(k)}, \hat{x}_i^{(\ell)}) &= \frac{1}{2} \int_0^1 \left| \hat{s}_j^{(k)}(f) - \hat{s}_i^{(\ell)}(f) \right| df \\ &= \frac{1}{2} \int_0^1 \left| s^{(k)}(f) + w^{(k)}(f) + e_j^{(k)}(f) \right. \\ &\quad \left. - (s^{(\ell)}(f) + w^{(\ell)}(f) + e_i^{(\ell)}(f)) \right| df \\ &\leq d(X^{(k)}, X^{(\ell)}) + \frac{1}{2} \int_0^1 |e_j^{(k)}(f)| df \\ &\quad + \frac{1}{2} \int_0^1 |e_i^{(\ell)}(f)| df \\ &\leq d(X^{(k)}, X^{(\ell)}) + \varepsilon. \end{aligned} \quad (6)$$

Reversing the triangle inequality leading to (6) it follows similarly that

$$d(\hat{x}_j^{(k)}, \hat{x}_i^{(\ell)}) \geq d(X^{(k)}, X^{(\ell)}) - \varepsilon \quad (8)$$

for all $k, \ell \in [L], i \in [n_\ell], j \in [n_k]$. Replacing $d(\hat{x}_i^{(\ell)}, \hat{x}_j^{(\ell)})$ on the RHS of (4) by the upper bound in (7) and using the lower bound in (8) on the LHS of (4), we find that (4) is implied by

$$\min_{k, \ell \in [L]: k \neq \ell} d(X^{(k)}, X^{(\ell)}) > 2\varepsilon. \quad (9)$$

We continue by upper-bounding ε . To this end, define $\mathbf{Q}_m \in \{0, 1\}^{M \times M}$ as $(\mathbf{Q}_m)_{u,v} = 1$, if $v - u = m$, and $(\mathbf{Q}_m)_{u,v} = 0$, otherwise, $\mathcal{M} := \{-M + 1, -M + 2, \dots, M - 1\}$, and $\mathbf{G}(f) := \sum_{m \in \mathcal{M}} g[m] \cos(2\pi f m) \mathbf{Q}_m$, i.e., $\mathbf{G}_{u,v}(f) = \mathbf{G}_{v,u}(f) = g[v - u] \cos(2\pi f(v - u))$. Now,

²available at <http://mocap.cs.cmu.edu>

with $\mathbf{x} \in \mathbb{R}^M$ the random vector containing the elements of $\hat{x}_i^{(\ell)}$, it holds for $m \in \mathcal{M}$ that

$$\hat{r}_i^{(\ell)}[m] = \frac{\mathbf{x}^T \mathbf{Q}_m \mathbf{x}}{M} \quad \text{and} \quad \tilde{r}^{(\ell)}[m] = \frac{\mathbb{E}[\mathbf{x}^T \mathbf{Q}_m \mathbf{x}]}{M - |m|}.$$

With these relations we have

$$\begin{aligned} \sup_{f \in [0,1]} |e_i^{(\ell)}(f)| &= \sup_{f \in [0,1]} |\hat{s}_i^{(\ell)}(f) - \tilde{s}^{(\ell)}(f)| \\ &= \sup_{f \in [0,1]} \left| \sum_{m \in \mathcal{M}} g[m] \hat{r}_i^{(\ell)}[m] e^{-i2\pi f m} - \sum_{m \in \mathbb{Z}} \tilde{r}^{(\ell)}[m] e^{-i2\pi f m} \right| \\ &= \sup_{f \in [0,1]} \left| \sum_{m \in \mathcal{M}} \frac{g[m]}{M} (\mathbf{x}^T \mathbf{Q}_m \mathbf{x} - \mathbb{E}[\mathbf{x}^T \mathbf{Q}_m \mathbf{x}]) e^{-i2\pi f m} \right. \\ &\quad \left. + \sum_{m \in \mathcal{M}} \frac{g[m]}{M} \underbrace{\mathbb{E}[\mathbf{x}^T \mathbf{Q}_m \mathbf{x}]}_{=g[m](1-\frac{|m|}{M})\tilde{r}^{(\ell)}[m]} e^{-i2\pi f m} - \sum_{m \in \mathbb{Z}} \tilde{r}^{(\ell)}[m] e^{-i2\pi f m} \right| \\ &= \sup_{f \in [0,1]} \left| \frac{1}{M} (\mathbf{x}^T \mathbf{G}(f) \mathbf{x} - \mathbb{E}[\mathbf{x}^T \mathbf{G}(f) \mathbf{x}]) \right. \\ &\quad \left. - \sum_{m \in \mathbb{Z}} h[m] \tilde{r}^{(\ell)}[m] e^{-i2\pi f m} \right| \end{aligned} \quad (10)$$

$$\leq \sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| + \sum_{m \in \mathbb{Z}} |h[m]| |r^{(\ell)}[m]| + \underbrace{|h[0]|}_{=0} \sigma^2 \quad (11)$$

$$= \sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| + \mu_{\max}, \quad (12)$$

where we used the fact that $g[m](\mathbf{x}^T \mathbf{Q}_m \mathbf{x} - \mathbb{E}[\mathbf{x}^T \mathbf{Q}_m \mathbf{x}])$ is a real-valued even sequence, employed the definition of $h[m]$ in (2) in the step leading from (10) to (11), and invoked $g[0] = 1$ (i.e., $h[0] = 0$) as well as $\tilde{r}^{(\ell)}[m] = r^{(\ell)}[m] + \sigma^2 \delta[m]$ to obtain the last inequality. It follows from (12) that $\varepsilon \leq \max_{\ell \in [L], i \in [n_\ell]} \sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| + \mu_{\max}$ and hence (3) implies (4) via (9) on the event

$$\mathcal{F} := \left\{ \max_{\ell \in [L], i \in [n_\ell]} \sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| < 4A(B + \sigma^2) \sqrt{\frac{2 \log M}{M}} \right\}.$$

It remains to lower-bound $P[\mathcal{F}]$, which will be accomplished by upper-bounding the tail probability of the random variables $\alpha_i^{(\ell)}(f)$. For fixed f , the distribution of $\alpha_i^{(\ell)}(f)$ obeys

$$\alpha_i^{(\ell)}(f) \sim \frac{1}{M} (\mathbf{y}^T \mathbf{C}^T \mathbf{G}(f) \mathbf{C} \mathbf{y} - \mathbb{E}[\mathbf{y}^T \mathbf{C}^T \mathbf{G}(f) \mathbf{C} \mathbf{y}]),$$

where the entries of \mathbf{y} are independent standard normal random variables and $\mathbf{C} = (\mathbf{R} + \sigma^2 \mathbf{I})^{1/2} \in \mathbb{R}^{M \times M}$ with $\mathbf{R}_{u,v} = r^{(\ell)}[v - u]$ the (Toeplitz) covariance matrix corresponding to M consecutive elements of $X^{(\ell)}$. Setting $\mathbf{B} := \mathbf{C}^T \mathbf{G}(f) \mathbf{C}$, we can establish a bound on the tail probability of $\alpha_i^{(\ell)}$ by invoking a well-known concentration inequality for quadratic forms in Gaussian random vectors [16, Lem. 1], namely

$$\begin{aligned} P[|\mathbf{y}^T \mathbf{B} \mathbf{y} - \mathbb{E}[\mathbf{y}^T \mathbf{B} \mathbf{y}]| \\ \geq \|\mathbf{B} + \mathbf{B}^T\|_F \sqrt{\delta} + 2\|\mathbf{B}\|_{2 \rightarrow 2} \delta] &\leq 2e^{-\delta}. \end{aligned} \quad (13)$$

Next, we note that $\|\mathbf{B} + \mathbf{B}^T\|_F \leq 2\|\mathbf{B}\|_F \leq 2\sqrt{M}\|\mathbf{B}\|_{2 \rightarrow 2}$ and $\|\mathbf{B}\|_{2 \rightarrow 2} \leq \|\mathbf{C}\|_{2 \rightarrow 2}^2 \|\mathbf{G}(f)\|_{2 \rightarrow 2} = \|\mathbf{R} + \sigma^2 \mathbf{I}\|_{2 \rightarrow 2} \|\mathbf{G}(f)\|_{2 \rightarrow 2} \leq A(B + \sigma^2)$, where we used that both \mathbf{R} and $\mathbf{G}(f)$ are symmetric Toeplitz matrices and hence, by [17, Lem. 4.1], $\|\mathbf{R}\|_{2 \rightarrow 2} \leq \sup_{f \in [0,1]} s^{(\ell)}(f) \leq B$ and $\|\mathbf{G}(f)\|_{2 \rightarrow 2} \leq \sup_{f' \in [0,1]} g(f + f') = \sup_{f' \in [0,1]} g(f') \leq A$ (here, the frequency shift by f due to the cos-factors in the elements of $\mathbf{G}(f)$ does not affect the supremum because $g(f)$ is 1-periodic). Now, setting $\delta = 2 \log(M)$ and using $\delta/M < \sqrt{\delta/M} < 1$, for $M \geq 1$, (13) yields

$$P \left[\sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| \geq 4A(B + \sigma^2) \sqrt{\frac{2 \log M}{M}} \right] \leq \frac{2}{M^2}.$$

Finally, it follows from a union bound argument that

$$\begin{aligned} P[\mathcal{F}] &\geq 1 - \sum_{\substack{\ell \in [L], \\ i \in [n_\ell]}} P \left[\sup_{f \in [0,1]} |\alpha_i^{(\ell)}(f)| \geq 4A(B + \sigma^2) \sqrt{\frac{2 \log M}{M}} \right] \\ &> 1 - \frac{2N}{M^2}. \end{aligned}$$

REFERENCES

- [1] Y. Wang, Z. Liu, and J.-C. Huang, "Multimedia content analysis – Using both audio and visual clues," *IEEE Signal Processing Magazine*, vol. 17, no. 6, pp. 12–36, 2000.
- [2] K. Kalpakis, D. Gada, and V. Puttagunta, "Distance measures for effective clustering of ARIMA time-series," in *Proc. IEEE International Conference on Data Mining (ICDM)*, 2001, pp. 273–280.
- [3] M. Corduas and D. Piccolo, "Time series clustering and classification by the autoregressive metric," *Computational Statistics & Data Analysis*, vol. 52, no. 4, pp. 1860–1872, 2008.
- [4] J. Boets, K. De Cock, B. De Moor, and M. Espinoza, "Clustering time series, subspace identification and cepstral distances," *Communications in Information & Systems*, vol. 5, no. 1, pp. 69–96, 2005.
- [5] J. Caiado, N. Crato, and D. Peña, "A periodogram-based metric for time series classification," *Computational Statistics & Data Analysis*, vol. 50, no. 10, pp. 2668–2684, 2006.
- [6] Y. Kakizawa, R. H. Shumway, and M. Taniguchi, "Discrimination and clustering for multivariate time series," *Journal of the American Statistical Association*, vol. 93, no. 441, pp. 328–340, 1998.
- [7] J. A. Vilar and S. Pértiga, "Discriminant and cluster analysis for Gaussian stationary processes: Local linear fitting approach," *Journal of Nonparametric Statistics*, vol. 16, no. 3-4, pp. 443–462, 2004.
- [8] D. Ryabko, "Clustering processes," in *Proc. 27th International Conference on Machine Learning (ICML)*, June 2010, pp. 919–926.
- [9] A. Khaleghi, D. Ryabko, J. Mary, and P. Preux, "Online clustering of processes," in *Proc. 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012, pp. 601–609.
- [10] Y. Xiong and D.-Y. Yeung, "Time series clustering with ARMA mixtures," *Pattern Recognition*, vol. 37, no. 8, pp. 1675–1689, 2004.
- [11] R. Heckel and H. Bölcskei, "Robust subspace clustering via thresholding," *arXiv:1307.4891*, 2014, submitted to *IEEE Transactions on Information Theory*.
- [12] I. Katsavounidis, C.-C. Jay Kuo, and Z. Zhang, "A new initialization technique for generalized Lloyd iteration," *IEEE Signal Processing Letters*, vol. 1, no. 10, pp. 144–146, 1994.
- [13] U. von Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, no. 4, pp. 395–416, 2007.
- [14] M. Soltanolkotabi, E. Elhamifar, and E. J. Candès, "Robust subspace clustering," *The Annals of Statistics*, vol. 42, no. 2, pp. 669–699, 2014.
- [15] L. Li and B. A. Prakash, "Time series clustering: Complex is simpler!" in *Proc. 28th International Conference on Machine Learning (ICML)*, 2011, pp. 185–192.
- [16] L. Demanet, P.-D. Létourneau, N. Boumal, H. Calandra, J. Chiu, and S. Snelson, "Matrix probing: A randomized preconditioner for the wave-equation Hessian," *Applied and Computational Harmonic Analysis*, vol. 32, no. 2, pp. 155–168, 2012.
- [17] R. M. Gray, "Toeplitz and circulant matrices: A review," *Foundations and Trends in Communications and Information Theory*, vol. 2, no. 3, pp. 155–239, 2006.